arXiv:1411.4411v3 [stat.ME] 3 Nov 2015

Ecological fallacy and covariates in the estimation of voters transitions

A. Forcina, Dipartimento di Economia, Finanza e Statistica,
University of Perugia, Italy

November 4, 2015

Abstract

In the estimation of voters transitions, additional insights are provided by a reformulation of the conditions for absence of ecological bias. With 2×2 tables, a graphical display, introduced to visualize individual and ecological data, simultaneously, indicates how to construct artificial instances of ecological bias of any complexity. Concerning the possibility of correcting ecological bias by modelling the effect of covariates, we argue that the crucial issue is whether voting decisions depend on the individual value of covariates or on their polling station averages. Evidence from the analysis of real and simulated data seem to support this conjecture.

Keywords. Ecological inference, effect of covariates, voting behaviour, logistic models, mixture models..

1 Introduction

After an election, political scientists, as well as politicians, may be interested in estimates of voters transitions to answer question like, for instance, (i) what proportion of voters remained faithful to the party they voted before, (ii) what is the proportion of protest voters who moved into abstention, (iii) in a runoff election, what is the proportion of strategic voters, that is those who did not abstain, even if the candidate they voted previously was excluded from the race. In certain contexts it might also be interesting to know how people voted depending on race, social class or education. All these questions concern how the voting decisions in a given election depend on the decision in a previous election (or on voters characteristics, like race and education); however, due to the special nature of electoral data, the true joint distribution between such pairs of variables cannot usually be observed and the only source of information can be extracted either from official electoral data aggregated at the level of polling stations or from sample surveys.

Though sample surveys are based on direct answers given by voters, their results may be heavily biased when the proportion of non respondents is non negligible; in addition, sample survey, usually, do not provide accurate estimates at the level of small areas. An advantage of ecological data is that, in several countries, official electoral data are freely available for each polling station; these have a low level of aggregation which makes it possible to compute ecological estimates of voters transitions for small areas. A comparison of the performance

of sample surveys relative to that of several ecological inference methods for estimating the association between race and propensity to register in a New Orleans majoral election was attempted by Liu (2007). He showed that, relative to the proportions of registered voters provided by the true individual data, most ecological inference methods did better than survey data in that specific context.

Robinson (1950) was perhaps the first to perform a formal investigation of the main source of ecological bias; however, partly because his derivations are based on the correlation coefficient, the most direct implications of his findings are not so easy to understand from his equations. In addition, it is difficult to agree with the spirit of his concluding remark: "The purpose of this paper will have been accomplished, however, if it prevents the future computation of meaningless correlations..". Subramanian et al (2009) have shown that the raw association in the individual level, which Robinson considered as an objective truth, is actually equivalent to the estimate provided by a rather simple logistic model which, for the particular data set used by Robinson, turns out to be inappropriate.

The literature on ecological bias is extensive and keeps growing, thus a thorough discussion in beyond the scope of this paper. As Wakefield (2004), Sec. 3.2, pointed out, there are features of the ecological fallacy which are specific to studies in epidemiology where, usually, one is interested in the causal effect of a quantitative variable X measuring exposure on a binary response Y, like having or not a given disease. In the simplest case, with a response function which is non linear (see, for instance Rchardson et al, 1987), the bias will depend on the heterogeneity of X within each local unit and could be corrected by combining aggregated data with samples of individual data (Salway and Wakefield, 2007; Prentice and Sheppard, 1995). Greenland and Morgenstern (1989) and Greenland and Robins (1994) show, by a set of suitably constructed examples, that even when X is categorical, the presence of covariates, acting as confounders or effect modifiers, may produce a kind of bias which is specific to aggregated data and more difficult to correct.

In this paper we derive an explicit characterization of the conditions for ecological bias in the context of voters transitions and, more generally, of two-way tables; this partly extends and generalizes similar results appeared in the literature (in particular Firebaugh, 1978; Piantadosi et al, 1988; Wakefield, 2004; Gelman et al, 2001). In short, estimates are going to be biased whenever the propensity to vote a given party within local units, is correlated with the distribution of votes in the previous election. When this happens, estimates will be invalid irrespective of the ecological inference method that is used; this suggests that, for instance, most of the heated debates between King (1999) and some of his opponents like (Cho, 1998) is not well posed.

We investigate whether bias can be removed by using the distribution of votes in the previous election as covariates; in particular, we consider two available methods of ecological inference on two-way tables where the effect of covariates on transition probabilities may be modelled on the logit scale: the ordinary least squares of King (see Rosen et al, 1991, Sec. 4) and a modified version of the Brown and Payne model (see Forcina et al, 2012). Concerning the underlying mechanism by which covariates may affect voting decisions, we consider two alternative scenarios: (i) covariates measured at the level of polling station act as macropoperties (see Firebaugh, 1978, Sec. II) having a direct effect on voting decisions or (ii) covariates measure the relative size of groups whose members have a different voting behaviour. While modelling the effect of covariates will usually correct bias in (i), it is likely to fail in (ii). We provide empirical evidence for these findings by analysing real and simulated

sets of individual data. The real data come from the city of Palermo where voters are classified according to sex, age and the decision to vote (or not to vote) at a primary election of the Democratic Party.

In Section 2 we derive an explicit condition for the ecological fallacy to arise and illustrate this result both with artificial and real world data analysed with three different methods of ecological inference. The properties of ecological inference with covariates is discussed in Section 3 where we characterize situation where modelling covariates may fail to correct for the ecological bias. Conclusions and additional discussions are presented in Section 4.

2 A new look into the ecological fallacy

2.1 Notations

Let Y denote the choice of a voter among the C options available in a given election and X the corresponding choice among the R options available in a previous (or simultaneous) election; X may also be a categorical variable which we expect to be associated with Y, like, for instance, the age group, the sex or the social class. Suppose we are interested in the association between X and Y within a relatively small area, like a town, which is divided into a collection of N polling stations.

Let n_{uij} denote the number of voters in polling station u with $X = i$ and $Y = j$; in the context of vote transitions, these are the values we would observe if we had access to individual data. Except in very special cases, these frequencies cannot be observed and we have to relay on the aggregated data: n_{ui+}, the number of voters with $X = i$ and n_{u+j}, those who voted option j in the new election. Let $f_{uij} = n_{uij}/n_{ui+}$ denote the proportion of voters who choose option j among those with $X = i$ in polling station u; let also $y_{uj} = n_{u+j}/n_{u++}$ and $x_{ui} = n_{ui+}/n_{u++}$ denote the marginal proportions in the X, Y table. Let $\underline{\mathrm{F}}_u$ denote the table with elements f_{uij} and $\bar{F}$ the table corresponding table of proportions after aggregation of local units:

$$\bar{f}_{ij} = \frac{n_{+ij}}{n_{+i+}} = \sum_u f_{uij} \frac{n_{ui+}}{\sum_u n_{ui+}}; \tag{1}$$

because the $\bar{f}_{ij}$ are a weighted averages of the f_{uij}, one could take as the target of ecological inference the association between X and Y in the overall table.

2.2 Artificial examples

It may be instructive to look at some artificial examples with 2×2 tables where there are only few polling stations, this shows the ecological fallacy at work in a way which is simple and striking at the same time. Any table of vote transitions may be reduced into a 2×2 table by selecting one voting option by row and collapsing the rest; similarly for the columns. A 2×2 table of proportions that sum to 1 by row may be plotted as a line going from $(0,\ f_{u12})$ to $(1,\ f_{u22})$ with slope $f_{u22} - f_{u12}$. Because

$$y_{uj} = (1 - x_{u2}) f_{u12} + x_{u2} f_{u22},$$

the point corresponding to the marginal proportions $(x_{u2},\ y_{u2})$, that is the ecological observation, lies on the same line and may be anywhere, being determined only by the value of x_{u2}. The table of overall proportions will correspond to the line going from $(0,\ \bar{f}_{12})$ to $(1,\ \bar{f}_{2,2})$.

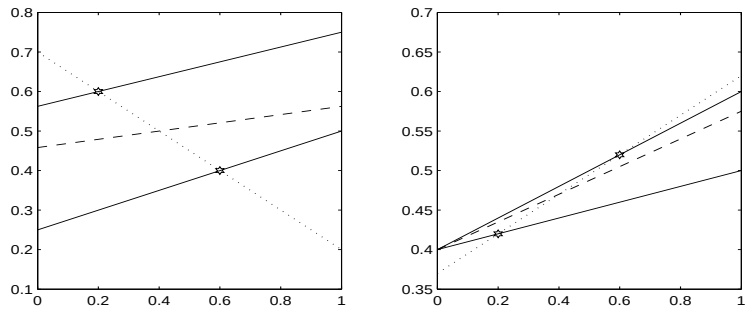

Figure 1: Two examples of ecological fallacy with two local units; solid lines join observations within each local unit, the dashed line represent the overall table, hexagram denotes ecological observations and the dotted line is the ecological regression

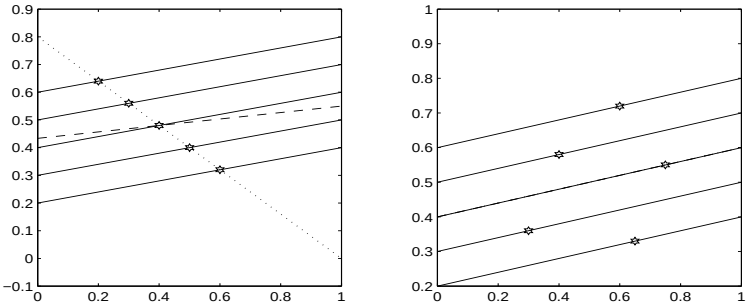

Figure 2: Two examples of ecological data with five local units; solid lines join observations within each local unit, the dashed line represent the overall table, hexagram denotes ecological observations and the dotted line is the ecological regression

The left panel of Figure 1 displays a situation where individual lines have similar positive slopes, different intercepts and an ecological regression line with negative slope while in the right panel individual lines have same intercepts, different slopes and an ecological regression line whose slope is greater than that of any individual line. Note that, in both cases, local units with a different value of x_{u2} have also different intercepts and/or different slopes.

To mimic an analysis of covariance model, in both examples in Figure 2 the lines joining the points within each local unit have constant slope but different intercepts. The left panel is an example where, though the ecological observations lie on a straight line, the slope of the ecological regression is negative while that in each local unit is positive and constant. This is obtained by generating data where the intercept f_{u12} decreases when x_{u2} increases. The data for the example displayed in the second panel of the same figure has been constructed by making the intercept f_{u12} vary as a quadratic function of x_{u2}. Here the x_{u2} points were chosen intentionally so that ecological bias in different directions cancel; a situation that will probably never arise in practice.

2.3 Ecological fallacy and linear regression

To discuss the general case where we have a collection of N polling stations with an $R \times C$ table of association, it is convenient to start with Goodman (1953) linear regression model. This is equivalent to assume that voters with $X = i$ choose option $Y = j$ with a probability π_{ij} which is constant across polling stations and that voters choose independently from each other. These assumptions imply that n_{u++} splits among the C options available in the second election according to a sum of R multinomial distributions. Starting from the accounting equation $y_{uj} = \sum x_{ui} f_{uij}$, with simple algebra (see the Supplementary material), we may write the regression equation as

$$y_{uj} = \pi_{1j} + \sum_{2}^{R} x_{ui}(\pi_{ij} - \pi_{1j}) + \epsilon_{uj},$$

where the error term $\epsilon_{uj} = \sum_1^R x_{ui}(f_{uij} - \pi_{ij})$. The condition for the least square estimates to be unbiased (Wooldridge, 2004, Sec. 2.2), is that the error term ϵ_{uj} is uncorrelated with the x_{ui} proportions, which are the explanatory variables in the regression model. Because this condition has to hold for all possible x_{iu}, in turn, it implies that we must have

$$E[f_{uij} - \pi_{ij} \mid x_{u1}, \ldots, x_{uR}) = 0. \tag{2}$$

The previous results may be summarized by saying that a necessary and sufficient condition for ecological regression to provide unbiased estimates is that the f_{uij} proportions are uncorrelated with the marginal proportions x_{ui}. Note, however, that the condition for the linear regression model to provide unbiased estimates are satisfied even if the f_{uij} vary at random or depend upon other variables, as long as these other variables are uncorrelated with the set of marginal proportions x_{ui}.

2.4 Ecological fallacy and the analysis of covariance

The analysis of covariance model has been used by Firebaugh (1978), Rchardson et al (1987), Piantadosi et al (1988) among others to study properties of ecological regression; in particular, Firebaugh stated the so called "$\bar{X}$-rule" which, translated within our context, is essentially equivalent to (2), with certain differences to be specified below. Both Firebaugh (1978) and Piantadosi et al (1988) define two regression coefficients for individual data, one based on a single local unit (assumed to be the same for all units) and the other on the overall data. Then they derive expressions showing how these coefficients are related with the ecological regression coefficient.

Relative to the context of vote transitions, the ordinary covariance model is too general and more restrictive at the same time. It is too general because it allows the marginal distribution of the covariates to vary arbitrarily from one local unit to the other, see for instance Figure 1 in Rchardson et al (1987); in the context of discrete data, covariates at individual level are dummy variables taking only value 0 or 1, as we saw in Sec. 2.2. Without this restriction, when we compute the average of Y conditionally on covariates for the overall individual data, we are also selecting specific subsets of local units; this is why the regression of Y on covariates in the overall individual data is a biased summary of the regression within local units. Instead, as indicated by (1), this cannot happen with categorical variables.

In our context, if we restrict for simplicity to the 2×2 case, when (2) is violated, $E(f_{u22} - f_{u12} \mid x_{u2})$, the slope of the regression line within local unit u, will usually depend on the local unit, while it is assumed to be constant in the analysis of covariance model which also assumes that, when ecological bias is present, $E(\epsilon_{u2} \mid x_{u2})$ is linear in x_{u2} which is the equivalent of $\bar{X}$ of Firebaugh in the bivariate case.

2.5 Ecological fallacy and logistic regression models

Though the model proposed by Brown and Payne (1986) did not attract much attention among political scientists, it is an improvement relative to ecological regression in several ways which we summarize below:

- the transition probabilities $p_{uij} = P(Y = j \mid X = i, U = u)$ are no longer assumed to be constant but may vary at random as in a Dirichlet distribution across polling stations around their expected value, say π_{uij};
- the transition probabilities π_{uij} are estimated on the logit scale, thus estimates always lie between 0 and 1; they may also depend on covariates on the same scale;
- the estimation method takes the covariance structure into account and thus, if the model assumptions are correct, its estimates will be more efficient;

see also Cleave et al (1995) for discussion and comparisons. A bayesian model with, essentially, the same likelihood was proposed by Greiner and Quinn (2009). A modification of the Brown and Payne model has been proposed by Forcina et al (2012) who assume that voters with the same $X = i$ living in the same polling station tend to cluster in smaller groups of random size such that voters who belong to the same cluster share the same probability of voting $Y = j \mid X = i$; these transition probabilities are assumed to vary again as in a Dirichlet distribution around their expected value. In practice, this model, which we will apply in the following, differs from the Brown and Payne model only for a component in the covariance matrix; readers interested in the technical details are referred to Forcina et al (2012)

G. King and coworkers have proposed several models of ecological inference; for simplicity, in the following we will consider the least squares version of their Multinomial-Dirichlet model (Rosen et al, 1991, Sect. 4). This model resembles the Brown and Payne model because both allow transition probabilities to vary at random across local units and transition probabilities are estimated on the logit scale. However, while the marginal frequencies of the votes at the second election is a convolution of Multinomial-Dirichelet distributions in Brown and Payne (1986), in Rosen et al (1991) this is assumed to be a single Multinomial with a vector of probabilities which is a convolution of Dirichlet distributions. This assumption was criticized by Wakefield (2004) but N. Best and King himself, in the discussion at the end of the same paper, argue in favour of this approach. In any case, when transition probabilities are estimated by OLS, the actual distinction is irrelevant because the covariance matrix is not taken into account.

Both the King and Brown and Payne models allow the transition probabilities to depend on covariates measured at the level of polling stations, a feature which will be considered in detail in Section 3. Without covariates, both models have an error term with a structure similar to the one in the ecological regression. It follows that, in order for the estimates to be consistent (see Wooldridge, 2004, Sec. 12.2), condition (2) must always be satisfied. In other

Table 1: Estimated proportions of people going to vote by sex and age in the Palermo data

	Individual data	Goodman	King	Brown-Payne
F	0.047	0.096	0.096	0.066
M	0.055	0.000	0.000	0.034
up to 45	0.045	0.000	0.000	0.114
over 45	0.055	0.091	0.091	0.001

words, if the condition is violated, we may expect that even the estimates provided by these more sophisticated models will be systematically different from the truth.

2.6 The Palermo data

In the city of Palermo, south Italy, after the 2012 Democratic Party primary election to select the candidate to run as major, individual data concerning sex, age and the choice to go to vote were made available for each of the 593 ordinary polling stations, excluding hospitals, prisons and other temporary locations. Table 1 gives the estimates of the proportion of voters by sex and age category in the overall individual data together with those obtained with the ecological inference methods presented above. Because the estimates from all ecological inference methods are substantially different from those resulting from individual data, it seems unlikely that this may be due to random variation. The theory discussed above suggests that this may be caused by the correlation between the propensity to vote and the marginal proportion of voters by age and sex in each polling station. A simple diagnostic which confirms this conjecture is given in Figure 3 which was constructed as follows: for the relation between sex and voting decision, we group polling stations according to the quantiles of the proportion of female voters and then plot the average proportion (y) of male and female voters for the polling stations within each pair of adjacent quantiles against the average proportion (x) of females voters. A similar plot is constructed by grouping polling stations on the basis of quantiles of the proportion of voters aged over 45 by plotting the average proportion of voters among those aged up to 45 and over 45 for the polling stations within each pair of adjacent quantiles. It emerges clearly that the proportion of voters for both males and females increases with the proportion of females and, for both age groups, with the proportion of voters aged over 45.

2.7 The 1930 US Census data

The data give the number of residents aged over 16 classified by race/nativity as "whites US born" (WUS), "Whites born aborad" (WBA) and "blacks" (BLK) and to whether they were "illiterate" or not. The results in Table 2 show the proportion of illiterate in the individual data and as estimated by the three ecological inference methods. Here all ecological inference methods seem to overestimate the proportion of illiterate among blacks.

Again, the theory above suggests that the bias should be due to correlation between the actual proportions of illiterate within each race/nativity group and the composition by race/nativity in different states. To examine this, in Figure 4 we grouped states on the basis of the quantiles of the proportion of WUS (left panel) and plotted the average proportion of illiterate against the average proportion of WUS within each quantile group. It emerges

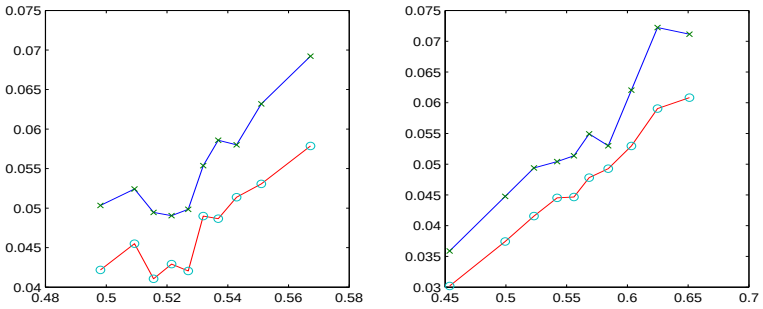

Figure 3: Quantile regressions: proportion of voters y by sex and proportion of females x in the left panel; proportion of votes by age and proportion of voters aged over 45 in the right panel; o stands for females or aged up to 45, $\times$ stands for males or aged over 45.

Table 2: Proportion of illiterate by race and nativity

	Individual data	Goodman	King	Brown-Payne
Whites US born	0.015	0.017	0.020	0.001
Whites non US born	0.098	0.027	0.014	0.082
Blacks	0.149	0.225	0.221	0.270

clearly that the proportions of illiterate within each group decrease when the proportion of white US born increase, though the effect is much smaller for whites and rather uneven for blacks. On the right panel, a similar plot was constructed by grouping states on the basis of quantiles of the proportion of WBA.

3 Ecological inference with covariates

When condition (2) is violated, the estimates produced by any method of ecological regression will be inconsistent, that is away from the truth even if the sample size (number of local units) was arbitrarily large. On the other hand, both King's ordinary least squares and Brown and Payne approaches, in their extended version, allow to model the effect of covariates, more precisely, we may define

$$\log \frac{\pi_{uij}}{\pi_{uiC}} = \alpha_{ij} + \sum_{k=1}^{K} \beta_{ijk} z_{uijk}, \tag{3}$$

where π_{uij} are the transition probabilities specific for local unit u, α_{ij} are intercept parameters on the logit scale, z_{uijk} are possible covariates measured at the level of polling stations and β_{ijk} are regression coefficients measuring the effect of the kth covariate on the transition probability in cell (i, j). Condition (2) tells us that we do not need to search among all possible covariates that may affect transition probabilities but simply restrict attention to the marginal proportions x_{ui}. Even then, when either or both R, C are large, there would be several different models to consider because each of the $R(C-1)$ entries may depend on any

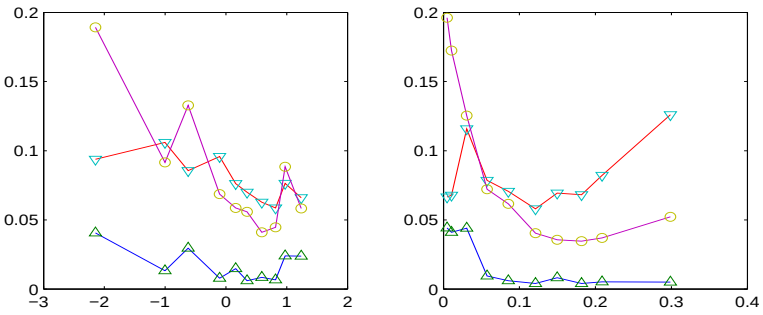

Figure 4: Quantile regression for the proportion of illiterate y with respect to the proportion (x) of whites US born on the left and, on the right, with respect to the proportion x of whites born abroad; $\triangle$ stands for whites US born, o for blacks and ∇ for other whites.

of the R marginal proportions; in addition, the z_{uijk} above could be non linear functions of the x_{ui}.

Suppose, for simplicity, that in (3) $K = 1$ so that the k index can be ignore, that all β_{ij} are 0 except when $i = j$ and that $z_{uii} = x_{ui}$, in other words, each diagonal entry of the table of transition probabilities depends on the corresponding marginal proportion. Such a model would hold if, for instance, the fact that a given party was quite strong in a polling station in the previous election has an influence on social relations which, in turn, affect individual decisions. In this case the x_{ui} proportions could be interpreted as macro-properties of local units in the sense of Firebaugh (1978). Suppose, instead, that there is an other variable, say V, for instance unemployment, which may affect voting decisions in both elections. Because voters will probably decide depending on their personal working condition, one would expect that the table of transition probabilities within a polling station is a mixture of two different tables, one for unemployed and one for the others with weights proportional to the relative size of the two groups. If so, modelling the effect of covariates on the logit scale may fail to correct the ecological bias.

When we fit an ecological inference model with covariates, the estimated proportion of voters with $Y = j$ among those with $X = i$ depends on the value of the covariates and thus is no longer constant across polling stations. Thus, in order to compare such estimates with the corresponding individual data, we need to reconstruct the estimated individual data within each polling station and then average across all polling stations.

In the following we will apply the approach described above to the Palermo data and to suitable sets of simulated individual data in an attempt to answer the question whether, by modelling the effect of covariates, we may correct ecological bias. For a detailed analysis of the US Census data, Firebaugh (1978) detected an interaction between being black and living in a state in the south; Subramanian et al (2009) also detected a strong effect of state regulations on racial segregation.

3.1 Applications to the Palermo data

To model the association between sex and propensity to vote in the Palermo data where both X and Y are binary, we simply assume that the logit of the probability to vote conditionally on

Table 3: Estimated proportions of people going to vote by sex and age, LI=logistic model on individual data, BP=Brown-Payne

Sex	LI	King	BP	Age	LI	King	BP
F	0.047	0.086	0.012	up to 45	0.045	0.097	0.115
M	0.055	0.012	0.095	over 45	0.055	0.014	0.000

sex is a linear function of the proportion of female voters within each polling station. As can be seen from Table 3, the estimates provided by both the King and the Brown-Payne models are far away from the truth, but in different directions. Though the association between the propensity to vote at the PD primaries and the proportion of female voters is clear from Figure 3, the proportion of females is unlikely to be a relevant property of polling stations that affects the tendency to vote directly. It seems more likely that the real determinant of voting at the PD primaries is affiliation to center-left parties who, at the election that took place two months later, got about 16.4% of the votes.

One possible explanation of the relation that causes ecological bias might be the following: polling stations with an higher proportion of voters affiliated to center-left parties are also those where voters are older and thus have an higher proportion of females. If so, the underlying mechanism might be a mixture model as described in Sec. 3 above. Though the logistic model assumed in both King and Brown-Payne is not a correct specification of the true underlying model, still it could have provided a reasonable approximation to the truth, thus it is a little surprising that results are so bad. To understand whether the failure is due to the inadequacy of the logistic approximation or to a feature which is related to the ecological nature of the data, we fitted a simple logistic model separately to the male and female voters by using the individual data and reconstructed the estimated proportions in each polling station and in the overall town. The fact that these estimates coincide up to the first three decimal digits with those based on the raw individual data indicates that the problem is not in the logistic approximation to the true underlying mechanism, but in the aggregation process itself.

The analysis for the relation between age and propensity to vote give similar results. When we try to correct the ecological bias due to the positive association between propensity to vote and proportion of voters aged over 45 by fitting ecological inference methods where the propensity to vote is allowed to depend on the appropriate covariate, estimates are not much better relative to those where propensity to vote vas assumed to vary at random across polling stations. Here again one might conjecture that the proportion of voters aged over 45 is not in itself a macro-property affecting the decision of voters directly; it seems more likely that voters decide depending on whether, individually, they are or are not affiliated to center-left parties and age is relevant only because older people are more likely to be affiliated.

Because individual data concerning the distribution of age, affiliation to center-left parties and voting at the PD primaries are not available, in the second example of the next section we will analyse data simulated under a mechanism which mimics what might have happened in the PD primaries and indeed, even if the sample size is very large, ecological inference methods adjusted for covariates are unable to correct the ecological bias under a mixture model-

Table 4: Estimates of transition probabilities in a context with three parties in both elections in an artificial example

	Individual data			Goodman		
	A	B	C	A	B	C
A	0.721	0.186	0.093	0.886	0.094	0.020
B	0.092	0.630	0.278	0.027	0.837	0.136
C	0.094	0.094	0.812	0.000	0.033	0.967
	King ols			Brown-Paynbe		
	A	B	C	A	B	C
A	0.721	0.192	0.087	0.728	0.186	0.086
B	0.097	0.628	0.274	0.094	0.630	0.275
C	0.091	0.090	0.819	0.087	0.093	0.819

3.2 Application to simulated data

It is possible to generate at random artificial electoral data from the modified Brown and Payne model. In short, the generation mechanism requires the user to set a basic table of transition probabilities which will hold when all covariates are equal to their averages and define a suitable logistic function for those entries which are to depend on covariates. The actual algorithm for generating the data is described in detail in the Supplementary Material. For each polling station the algorithm will produce a table of individual data with the joint distribution of voters according to the options available in the two elections. Having individual and aggregated data at the same time, we may check whether the results of ecological inference methods are consistent. Because the purpose is not to compare the efficiency of different estimation methods but simply to check whether consistent estimates may be obtained, for each example, we generated a single data set with 20,000 polling stations.

The first data set was generated from a 3×3 table of transition probabilities where the values on the main diagonal may be interpreted as corresponding to faithful voters and are, thus, the largest entries within each row. We also assume that the entries on the main diagonal increase with the proportion in the corresponding row margin as in a logistic function. It can be easily seen from Table 4 that both King's OLS and the modified Brown-Payne model provide estimates which are very close to the truth. On the other hand, the Goodman model, which does not account for the effect of covariates, clearly overestimates the transition probabilities along the main diagonal.

The second data set has $R = 2$, $C = 2$ and was generated under the following assumptions:

1. voters may be of two different types depending on a binary variable V, each type has its own matrix of transitions;

2. the proportion of voters with $V = 1$ vary across polling stations and is correlated with the decision X in the previous election;

3. conditionally on $V = v$ voters are distributed with respect to Y as in the modified Brown-Payne model.

To fix ideas, suppose that $X = 1$ means female, $Y = 1$ means vote in the primary election and

Table 5: True and estimated probability of voting: I=individual data, L=logistic approximation, G=Goodman regression, K=King OLS, BP=modified Brown-Payne

	True		Estimated				
Sex	$V=0$	$V=1$	I	L	G	K	BP
M	0.005	0.402	0.087	0.0874	0.162	0.019	0.006
F	0.001	0.400	0.078	0.0783	0.003	0.147	0.160

$V = 1$ for voters affiliated to a center-left party. Table 5 gives the true proportion of voters with $Y = 1$ in the generating process, the corresponding proportions in the overall individual data and as estimated by the different ecological inference methods.

The results in Table 5 clearly indicate that by modelling the single covariate on the logit scale neither King OLS nor the modified Brown-Payne approach can correct the ecological bias. Because in these data the proportions of those who voted at the primaries, f_{ui2}, is correlated with x_{u2}, the proportion of females, (2) is violated. Though the logistic model assumed in both the King and Brown-Payne models is not a correct specification of the true generating mechanism, it still provides a very close approximation to the mixture model if we had access to the individual data. Thus the bias must be due to the ecological nature of the data.

4 Discussion and conclusions

This paper is an investigation into the bias which may affect ecological inference in the context of estimation of voters transitions when studying the association between voters decisions in two close in time elections. By exploiting basic results concerning unbiasedness and consistency of estimates in linear and non linear regression, we argue that, when certain conditions are violated, any methods of ecological inference will be inconsistent. In the light of these results, the crucial issue is not which ecological inference method is better because, as our applications indicate, the bias is going to affect any formal method of inference. This is in contrast with the debate within the sociological literature, see for example Seligson (2002), where the claim that certain methods can avoid ecological bias are taken seriously. This is partly a consequence of the fact that King (1997) entitled his paper "A solution to the ecological inference problem"; in certain section of the social science community this was taken by some as meaning that, at last, the problem of bias in ecological studies had been solved.

Concerning the possibility of modelling the effect of covariates, our results indicate that one does not need to search among all possible available covariates associated with local units but simply on the marginal proportions of voters who selected a given option in a previous election or with a given characteristic like sex, age or education. Unfortunately, unless individual data are available, as in the examples considered in this paper, there is no diagnostic for checking whether the condition for obtaining consistent estimates from aggregated data are satisfied. However, one could fit an ecological inference method which allows certain transition probabilities to depend on a specific covariate and then test whether this effect is significant. Clearly this procedure of try and test may be rather time consuming.

We also show that bias correction by modelling the effect of covariates may also fail for reasons related to the nature of the covariates. When covariates characterize certain general

features of a local unit, like, for instance, the fact that a certain party is more or less strong, because of social interactions, this is likely to affect certain categories of voters who live in the area. In other contexts, the behaviour of voters may be affected by their individual characteristics, like for instance, being unemployed or being affiliated to a given party, in the terminology Firebaugh (1978) these are "micro-properties". We argue that an appropriate model in such contexts is a mixture of two (or more) tables of transition probabilities with weights given by the proportion of voters with a given micro-property. Our results based on observed and simulated individual data indicate that using the proportion of voters as a covariate on the logit scale may fail completely to correct for ecological bias. In principle, one could think of ecological inference methods designed to handle mixtures where the weights are known. A method based on assuming that polling stations may be clustered according to different latent types, each with its own table of transition probabilities, has recently been proposed by Puig and Ginebra (2014).

In several papers (for example King, 1997) King claims that his approach is able to combine the method of bounds with the regression approach, a claim which seems to be based on the fact that the estimates in his bayesian models are computed conditionally on the observed margins. Even frequentist estimates may be forced to satisfy the marginal constraints by computing an estimated table of counts which has the same odds ratio as in the estimated table of transition probabilities and the same row and columns margins as observed in each polling station. Though this adjustment, which has been applied in all the examples presented in this paper, may produce slightly better estimates, it will not correct for ecological bias.

Acknowledgments

The author would like to thank two anonymous referees for comments and suggestions, V.. Tomaselli (Univ. of Catania) for providing the Palermo data, D.J. Greiner and G. King for discussion, S.V. Subramanian (Harward School of P.H.) for providing the Robinson's data and useful interactions.

References

Brown, P., Payne, C. (1986). Aggregate data, ecological regression and voting transitions. *Journal of the American Statistical Association*, 81, 453–460.

Cho, W.T.K. (1998). If the assumption fits: A comment on the King ecologic inference solution. *Political Analysis*, 7, 143-163.

Cleave, N., Brown, P.J., Payne, C.D. (1995). Evaluation of methods of ecological inference. *Journ. Royal Statist. Soc, A*, 158, 55–72.

Freedman, D.A., Klein, S.P., Ostland, M. and Roberts, M.R. (1998) On "solutions" to the ecological inference problem *Journal of the American Statistical Association*, **98**, 1518-1522.

Firebaugh, G. (1978). A rule for inferringindividual level relationships from aggregate data. *American Sociological Review*, 43, 557–572.

Forcina, A., Gnaldi, M. and Bracalente, B. (2012). A revised Brown and Payne model of voting behaviour applied to the 2009 elections in Italy. *Statistical Methods and Applications*, 21, 109–119.

Freedman, D. A., Ostland, M., Roberts, M. R., and Klein, S. P. (1999). Reply to King. *Journal of the American Statistical Association*, 94, 355-357.

Gelman, A., Park, D.K., Ansolabehere, S., Price, P.N., Minnite, L.C. (2001). Models, assumptions and model checking in ecological regressions. *Journ. Royal Ststist. Soc., A*, 164, 101–118.

Goodman, L. A. (1953). Ecological regression and the behaviour of individuals. *American Sociological Review*, 18, 269–274.

Greenland, S. and Morgenstern, H. (1989). Ecological bias, confounding, and effect modification. *Intern. J. of Epidemiology*, 18, 269–274.

Greenland, S. and Morgenstern, H. (1989). Ecological bias, confounding, and effect modification. *Intern. J. of Epidemiology*, 18, 269–274.

Greenland, S. and Robins, J. (1994). Ecological studies–Biases, misconseptions and counterexamples. *Amer. J. of Epidemiology*, 139, 747–760.

Greiner, D.J. and Quinn, K.M. (2009). R x C ecological inference: bounds, correlations, flexibility and transparency of assumptions. *J. Royal Statist. Soc. A*, 172, 67–81.

King, G. (1997). *A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data*. Princeton University Press, Princeton, NJ.

King, G. (1999). The future of ecological inference. Journal of the American Statistical Association, 94, 352-354.

Liu, B. (2007). EI extended model and the fear of ecological fallacy. *Sociological Methods and Research*, 20, 1–23.

Piantadosi, S., Byar, D.P., Green, S.B. (1987). The ecological fallacy. *Amer. J. of Epidemiology*, 127, 893–904.

Prentice, R.L. and Sheppard, L. (1995). Aggregate data studies of disease risk factors. *Biometrika*, 82, 113–125.

Puig, X. and Ginebra, J. (2014). A cluster analysis of vote transitions. *Computational Statistics and Data Analysis*, 70, 328–344.

Richardson, S., Stüker, I., Hémon, D. (1987). Comparison of relative risks obtained in ecological and individual studies: some methodological considerations. *Intern. J. of Epidemiology*, 16, 111–120.

Robinson, W.S. (1950). Ecological correlations and the behaviour of individuals. *Amer. Sociological Rev.*, 15, 351–357.

Rosen, O., Jiang, W., King, G. and Tanner, M.A. (2001). Bayesian and frequentist inference for ecological inference: the RxR case. *Statistica Neerlandica*, **55**, 134–156.

Russo, L. (2014). Estimating floating voters: a comparison between the ecological inference and the survey methods. *Quality and Quantity*, 48, 1667-1683.

Seligson, M.A. (2002). The renaissance of political culture or the renaissance of the ecological fallacy ? *Comparative Politics*, 34, 273–292.

Salway, R., Wakefield, J. (2007). Sources of bias in ecological studies of non-rare events. *Environmental and Ecological Statistics*, 12, 321–347.

Subramanian, SV., Jones, K., Kaddour, A. and Kreiger, N. (2009). Rivisiting Robinson: the perils of individualistic and ecologic fallacy. *Intern. Journ. of Epidemiology*, 38, 342–360.

Wakefield, J. (2004). Ecological inference for 2×2 tables. *Journal of the Royal Statist. Soc. A*, 167, 1–42.

Wooldridge, J.M. (2004). *Econometric analysis of cross section and panel data*. The MIT Press, Cambridge, Massachusetts.